

UTILITY PATENT APPLICATION TRANSMITTAL (Small Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
ADT0001-US

Total Pages in this Submission

TO THE ASSISTANT COMMISSIONER FOR PATENTS

Box Patent Application
Washington, D.C. 20231

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

SYSTEM AND METHOD FOR CREATING, DISTRIBUTING AND MANAGING ARTIFICIAL AGENTS

and invented by:

CHRISTOPHER R. STEPHENS AND HENRI WELBROECK

If a **CONTINUATION APPLICATION**, check appropriate box and supply the requisite information:

☒ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☒ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Enclosed are:

Application Elements

1. ☒ Filing fee as calculated and transmitted as described below
2. ☒ Specification having 53 pages and including the following:
 - a. ☒ Descriptive Title of the Invention
 - b. ☐ Cross References to Related Applications (if applicable)
 - c. ☐ Statement Regarding Federally-sponsored Research/Development (if applicable)
 - d. ☐ Reference to Microfiche Appendix (if applicable)
 - e. ☒ Background of the Invention
 - f. ☒ Brief Summary of the Invention
 - g. ☒ Brief Description of the Drawings (if drawings filed)
 - h. ☒ Detailed Description
 - i. ☒ Claim(s) as Classified Below
 - j. ☒ Abstract of the Disclosure

**VERIFIED STATEMENT BY A NON-INVENTOR SUPPORTING
A CLAIM BY ANOTHER FOR SMALL ENTITY STATUS**

Docket No.
ADT0001-US

Serial No.
NEW APPLICATION

Filing Date
HEREWITH

Patent No.
-

Issue Date
-

Applicant/ **CHRISTOPHER R. STEPHENS AND HENRI WELBROECK**
Patentee:

Invention: **SYSTEM AND METHOD FOR CREATING, DISTRIBUTING AND MANAGING ARTIFICIAL AGENTS**

I hereby declare that I am making this verified statement to support a claim by

for small entity status for purposes of paying reduced fees to the United States Patent and Trademark Office, regarding the invention described in:

- ☒ the specification filed herewith with title as listed above.
☐ the application identified above.
☐ the patent identified above.

I hereby declare that I would qualify as an independent inventor as defined in 37 CFR 1.9(c) for purposes of paying fees to the United States Patent and Trademark Office, if I had made the above identified invention.

I have not assigned, granted, conveyed or licensed and am under no obligation under contract or law to assign, grant, convey or license, any rights in the invention to any person who would not qualify as an independent inventor under 37 CFR 1.9(c) if that person had made the invention, or to any concern which would not qualify as a small business concern under 37 CFR 1.9(d) or a nonprofit organization under 37 CFR 1.9(e). Note: Separate verified statements are required from each person, concern or organization having rights to the invention averring to their status as small entities. (37 CFR 1.27)

Each person, concern or organization to which I have assigned, granted, conveyed, or licensed or am under an obligation under contract or law to assign, grant, convey, or license any rights in the invention is listed below:

- ☒ no such person, concern, or organization exists.
☐ each such person, concern or organization is listed below.

FULL NAME
ADDRESS

☐ Individual ☐ Small Business Concern ☐ Nonprofit Organization

FULL NAME
ADDRESS

☐ Individual ☐ Small Business Concern ☐ Nonprofit Organization

FULL NAME
ADDRESS

☐ Individual ☐ Small Business Concern ☐ Nonprofit Organization

FULL NAME
ADDRESS

☐ Individual ☐ Small Business Concern ☐ Nonprofit Organization

SYSTEM AND METHOD FOR CREATING, DISTRIBUTING AND MANAGING ARTIFICIAL AGENTS

FIELD OF THE INVENTION

The present invention relates to intelligent artificial agents and more particularly to systems and methods for the development, distribution and management of intelligent artificial agents.

BACKGROUND OF THE INVENTION

Technological progress has allowed machines to replace humans in a growing range of repetitive tasks, improving the efficiency of industrial processes and raising productivity in a dramatic way. This improvement however has not been fully matched in the service sector, in part because many of the tasks carried out in the services industry are perceived to require intelligent reasoning and skills that are not easily emulated by machines.

This is particularly true in finance where fundamental analysis has been a human prerogative and traditional technical analysis has also relied on humans to interpret its results. There are many technical analysis tools available, but these are often relatively unsophisticated systems that are viewed as “toys” by more quantitatively minded analysts. While there are comparatively better systems, such as Metastock, available from Equis International, Salt Lake City, Utah, that provide a large amount of potentially useful information, these systems are not true prediction systems since they rely, ultimately, on heuristic rules that cannot have global validity over time in a constantly changing financial market.

A large part of finance is associated, for obvious reasons, with trying to predict what will happen in the future. Both fundamental and technical analysis are concerned with the question of prediction but approach it from completely different points of view. Traditional economic theory (the Efficient markets hypothesis) states that markets are by definition unpredictable. However, there is

intelligent search of databases is carried out. In fact many such systems are no more sophisticated than standard web search engines.

Thus, there remains a need for a different approach to development, distribution and management (or use) of intelligent artificial agents, and particularly those that are applicable to dynamic systems such as stock markets where a crucial aspect of the agent system must be its ability to adapt to changing conditions and where prediction plays a fundamental role.

There are clear opportunities for using intelligent artificial agents as analytical tools, and particularly so in finance. In the Internet sector alone (i.e., online trading), as of 1999 there were more than 5.4 million active online brokerage accounts managing assets of nearly \$500 billion. This amount is expected to rise to \$3 trillion by 2003. Discount brokers are presently reaching the limit of how efficient they can make their operations and still make a profit in the midst of intense competition and falling transaction fees. Accordingly, there is an increasing tendency to add on new services that give substantial added value. In particular, investors continue to seek relevant and timely information that can give them a competitive edge in their trading decisions but at a price that does not excessively impact the overall cost of trading.

While the glory days of the full service broker will likely soon disappear, the need for advice on equity trading and other elements of financial assessment will not. Accordingly, there is a need for affordable advice in terms of both fundamental and technical analysis on equities and other financial instruments as well as a constant need for assessment of risk, tax positions, portfolio management, etc. Indeed, the potential market for timely and affordable advice is enormous. A conservative market penetration estimate of 1% of managed online assets by 2003 and total commissions and fees equivalent to 0.5% of assets managed generates revenues of \$150 million. Artificial agent systems are capable of offering superior

More specifically, the present invention can be considered a consulting system that is based primarily on the decision-making abilities of artificial agents. The system includes a graphical user interface to interact with investors and aide in selecting the appropriate artificial agents. The artificial agents of the present invention replace, in many respects, human agents. Importantly, the artificial agents provide superior trading strategies, which are based on the detection of statistical biases in financial time series.

A constantly changing global marketplace demands prediction tools of the highest quality. Today's demanding financial professionals and investors need the power of improved technology. The present invention offers them this power with an artificial agent-based prediction system, wherein each agent is created from an adaptive system and is, in its turn, an adaptive system. More specifically, the present invention includes the artificial agent factory where agents are created. In the factory, agents exist in a competitive "Darwinian" environment where an ecology of agent strategies compete. Survival of the fittest in this context means survival of the most predictable. Each agent that leaves the factory includes a quantitative predictability value which is a direct measure of the agent's "fitness" and which can be used by a user to help choose which particular agent or agents to employ.

With respect to artificial agent distribution, the leasing system of the present invention makes it possible to provide pre-trained, computationally powerful prediction agents adapted to current market conditions without placing an excessive burden on the user's own resources. Through this innovative solution years of intensive research can be made available to financial decision makers worldwide. Preferably, the artificial agents provide advice on buying and selling financial instruments such as equities.

The artificial agents in accordance with the present invention are preferably trained to recognize emergent price structures (i.e., "predictability bubbles"), design

- an agent editor tool to allow a user to create an agent with the characteristics of his choice.

Within the finance arena, similarly-styled intelligent, agent-based techniques are applicable to other output variables (interest rate derivatives, currency derivatives, etc.) and other inputs including fundamental analysis data, news clippings, etc. The more this type of information is quantifiable, the more it can be used to augment the decision-making skills of the artificial agents of the present invention.

There is, of course, a need for improved artificial agent technologies in other areas. Accordingly, the agent-based tools of the present invention are also applicable to disciplines as diverse as the electricity supply industry and the insurance industry, where a search for and exploitation of predictability bubbles may be implemented.

Agent based modeling has drawn increased attention in recent years, in part for its value in helping to understand the behavior of complex adaptive systems, such as financial markets, and learn to forecast their evolution. For example, in portfolio management or day trading, artificial agents offer superior trading strategies and analysis when compared to their human counterparts.

The artificial agents of the present invention preferably include intelligence derived from a predictability value which itself is preferably further based on neural network information.

With respect to predictability, it is legion that “past performance is no guarantee of future returns.” However, a predictability indicator provided by the present invention has an objective criterion that links the past and future and directly measures an agent’s expectation of future success. The predictability

- b) decides how much money should be assigned to this consensus decision in light of the fund's present position in this instrument and in other instruments; and
- c) approves or rejects the final trade according to user-specified criteria relating to the fund's guidelines such as maximum exposure allowed per instrument type, or other known criteria.

The fund manager can decide to authorize this agent for automatic execution, in which case the recommended final trade is automatically passed on to an order execution agent. The order execution agent submits the order to an electronic trading system, manages the dialogue with this trading system and reports when the order is pending match or done. Reporting involves a message to the user (in this case a manager) and an automatic update of the user's portfolio database to reflect the new position.

Market Maker Tool

This tool seeks to confront the "adverse information cost" incurred by market makers when they do not immediately become aware of new information and respond to it by moving their quotes. This information can be technical in nature (e.g., price patterns that are deemed to be relevant by day-traders), news clippings, movements in the market makers montage, or information on trading decisions by large institutions, for example. Researchers estimate that about half of the quoted spread is lost to the adverse information risk. Much of this adverse information is in fact indirectly available to the market maker, to the extent that one can filter it out of the vast amount of incoming information describing everything from the market maker montage to customer order flow. This filtering task is similar in essence to the task of screening agents in the agent factory for the investor tool.

The stream of incoming orders at a given market maker provides valuable information which can be used to anticipate short-term price fluctuations and warn

Figure 3 is a simplified three-dimensional depiction of a landscape reconstruction.

Figure 4 is an exemplary curriculum vitae (CV) for an artificial agent in accordance with the present invention.

Figure 5 illustrates a stock manager user interface screen in accordance with the present invention.

Figure 6 illustrates an agent manager user interface screen in accordance with the present invention.

Figure 7 illustrates a ticker user interface screen in accordance with the present invention.

Figure 8 illustrates an agent round table user interface screen in accordance with the present invention.

Figure 9 illustrates the stock manager user interface screen including an agent recommendation window in accordance with the present invention.

Figure 10 illustrates graphical analysis user interface screens in accordance with the present invention.

Figure 11 is a depiction of an agent editor tool user interface screen in accordance with the present invention.

Figure 12 is an example of the possible scenarios produced by the agent editor tool in accordance with the present invention.

Figure 13 is a flowchart illustrating how it is determined that an artificial agent should be retired for retraining in accordance with the present invention.

DETAILED DESCRIPTION OF THE INVENTION

The present invention relates to intelligent artificial agents and more particularly to systems and methods for the development, distribution and management of intelligent artificial agents.

The present invention is designed for systems that are in a constant state of flux, or, more particularly, are adaptive. Accordingly, the present invention is particularly suited to financial markets, which respond to the actions of various market participants. It has been found that, with respect to stocks for example, structures continuously emerge in price charts as diverse trading strategies interact to produce coherent movements. Representing the combined effects of all these trading activities, these structures are an example of a phenomenon known in physics as “emergence of collective excitations.” In finance, researchers refer to emerging structures as “predictability bubbles.”

Even the best financial traders eventually lose step when markets are adapting quickly. However, an intelligent artificial agent in accordance with the present invention can recognize predictability bubbles when they arise, track their evolution and, significantly, anticipate their end. Thus, the intelligent artificial agents (hereinafter “artificial agents” or “agents”) of the present invention operate to identify price predictability patterns in equity markets and provide specific recommendations to a user as to how best to exploit that pattern. Each agent preferably is self-monitoring in that the agent tracks the evolution of a predictability bubble and anticipates its (the agent’s) own future performance.

Figure 1 shows a general schematic diagram of a preferred embodiment of the present invention, which, in this case, is directed specifically to a financial services implementation. However, it should be evident to those skilled in the art of artificial agents that the artificial agent development, distribution and management methods and systems discussed herein are equally applicable to other disciplines

various recommendations (shown via dotted line in Figure 1) and automatically execute the recommended action. When an automated trade clearing system is not implemented, a user preferably receives the aforementioned recommendation and uses it exclusively, or as part of the basis for his ultimate trading decision making.

Self-monitoring artificial agent

A self-monitoring agent in accordance with the present invention is an algorithm that preferably takes as input a time series, such as a price series for a given financial instrument, and executes a sequence of operations whose final result is a (set of) recommendation(s) on appropriate actions to be taken by the user and a measure of the expected quality of these recommendations, i.e., the expected future performance of the agent. In the preferred embodiment of the present invention the output is a trading strategy that produces recommendations for trading that instrument, and the quality of the recommendations is equal to the profits that the agent expects will be earned by applying this trading strategy. In this case the future performance is directly related to the degree of predictability of the time series being considered and the future performance estimate is called “predictability” (see below for a discussion of the latter). The agent algorithm is created in the agent factory (discussed below) and preferably comprises two distinct algorithms: an A-brain and a B-brain. The former searches through a large library of technical templates looking for one of high predictability, while the latter relies on a neural network to modify, when necessary, the result of the A-brain.

The technical template preferably is a technical trading strategy itself such as a moving average, MACD, etc., which are familiar to those skilled in the art. This A-brain strategy gives rise to a yield curve which is then fed to the B-brain which comes from a library of neural networks with, for example, a self-organized critical learning algorithm or a conventional backpropagation algorithm (many of which are commercially available) and well-known to those skilled in the art.

Preferably, the purpose of the A-brain is to generate yield curves that have desirable properties, such as low volatility and high yield per unit time. The purpose of the B-brain is to detect changes in market tendency that would invalidate the A-brain strategy, i.e., to identify when a predictability bubble associated with a given strategy is about to end, and to amend the A-brain strategy if and when necessary. Thus, the A-brain strategy might yield a buy signal but be overruled by the B-brain which may negate or overturn the buy signal. The A-brain and B-brain together comprise an agent and yield a trading strategy and a subsequent yield curve that may be described in terms of a set of parameters including, for example, timescale(s) of the technical template used, predictability, aggressiveness, the duration in time units of the agent's strategy (i.e., the duration of its past trading), the agent's precision (i.e., the percentage of correct decisions in past trading), the total number of decisions taken in past trading, the total profit of the agent in past trading, force threshold (see below), type of B-brain used, etc. Also relevant are several "geometrical" variables associated with an agent's "trading canal". The latter is defined by taking the agent's yield curve for its past trading and finding a pair of parallel lines that give a best envelope to this curve (see for example the Value curve in Figure 5). Parameters preferably include the length, slope, initial and final width of the canal. Preferably, all or a subset of these parameter values encode the agent and, therefore, its complete trading strategy such that when the agent is transmitted via internet or other means of communication to a user, the values act as input parameters to a set of executable programs which recreate the agent and therefore trading strategy and recommendations associated with those parameters.

Predictability and Landscape Reconstruction

In accordance with the present invention, each agent is preferably self-monitoring such that it is capable of anticipating the evolution of its own

performance by analyzing the strength of the predictability bubble on which its strategies are based. To implement this functionality, the present invention includes a method for measuring predictability. Predictability is based on mutual-information-based reconstruction of a multivariate fitness landscape which evaluates an agent's ability to generalize into the future. This landscape also plays an important role in the evolution dynamics of agents at agent factory 115, as will be explained in more detail later herein.

In the context of financial markets, landscape reconstruction is the process of analyzing past market data and determining a predictability value for each agent or technical template (see below for more discussion of technical templates). More specifically, the purpose of this landscape reconstruction is to establish the correlation between an agent's past performance and its expected future performance. "Predictability" is by definition the statistical expected value of a target variable that measures future performance, such as the agent's profits over the next 10 days, the agent's profits until its next trade or some other performance measure. The predictability of a technical template is measured by creating an agent that applies a simple technical trading rule, such as "buy when the price is above the moving average", and measuring the predictability of that agent. To reconstruct the predictability landscape, one performs a large number of experimental agent runs using historical data. In each agent run one executes the following steps:

- 1) Choose an agent or create one from a template-based technical trading rule, e.g. 10 day moving average, Fourier with period 30 days, MACD, or any other of a number of well-known technical trading rules.
- 2) Compute the earnings (yield) curve that this agent would have produced had he followed this trading strategy over a particular trading window (e.g. data lines numbered 400 to 500 in an historical database where each data line may represent a day, an hour or a market tick).

Agent factory 115 preferably maintains two pools of agents: active agents that are made available to users and reserve agents that preferably are used to replace active ones when they are retired as a result of changing market conditions. The main functions fall in two groups: management of active or live agents, and production of new reserve agents.

First group: managing the pool of live agents

“Data” reads new daily market price data from a financial data provider (market data 120) and makes this information available to the other factory functions.

“Run agents” updates each existing agent’s recommendations in light of the new data and retires agents that have a negative predictability value.

“Refresh” looks for agents that have been marked as “retired” due to changing market conditions, places these in a “cemetery” database and replaces them with reserve agents (previously generated) that have a high predictability value and trading strategies that are diverse to those currently adopted by existing active agents, usually because they are using different prediction horizons or templates.

“Trades” looks for agent consensus decisions using various non-linear consensus rules that identify teams of two or more agents that agree to recommend a trade. For example, if there exists for a given stock an agent with a short term horizon and positive PFI (a so-called “market timing” agent) together with a medium to long term horizon agent and both with predictability above a certain threshold and there exists no similar agent with a contrary opinion, then a position will be opened if both agents are in agreement. Trades also reviews trading positions recommended by teams of agents in previous runs of the factory process and decides if such positions should be closed based on the new data.

“A-brain” preferably creates a preliminary artificial agent that uses the technical trading rule selected and generates a preprocessed signal (chart) available for further processing.

“B-brain” preferably iterates a process that implements a pre-trained neural network from a library of candidates and completes the agent’s training by letting it learn how to interpret the preprocessed signal generated by “A-brain”; for each choice of neural network it creates an agent and analyzes its predictability. The most predictable agent from these trials is kept in a database to be later “packaged” and sent to the ultimate user. Among the B-brains in this library are some candidates that apply strategies that would be recognized as common technical trading rules by those skilled in the art, and preferably other candidates that apply more complex interpretations of the pre-processed signals from the A-brain, including contextual learning and backpropagation neural networks that are also known to those skilled in the art.

In a preferred embodiment of the invention, each agent is provided with a Curriculum Vitae (CV), as shown in Figure 4, with which the user may evaluate the characteristics of the agents, judge its performance and select the agent most suited to his needs. Included are basic data such as the agent name, which technical template it uses, its horizon, its neural network training algorithm, the date on which the agent graduated from the agent factory and the agent’s aggressiveness, i.e., how often, on average, it recommends a trade every 100 days, or other fixed period of time.

The second part of the CV relates to the agent’s performance during its training in agent factory 115 and consists of the agent’s precision, i.e., of all the recommendations it made what percentage was correct, the actual number of decisions taken and the actual number of correct decisions. Correct in this context signifies that the price rose when the agent recommended to buy or the price fell when it recommended to sell. Also preferably provided is the average percentage

a ticker (Figure 7);
an agent round table (Figure 8);
an agent recommendation (Figure 9); and
graphical analysis (Figure 10).

As will be appreciated by those skilled in the art, some of the data depicted on one screen might also appear in other screens. This facilitates the use of management system 125 and provides reassurance to a user that the underlying data is consistent. It is noted that Figures 5-10 are based on selected stocks listed on the Mexican stock exchange and artificial agents 110 for those stocks.

Figure 5 depicts a stock manager user interface screen in which stocks for which agents are available are listed under the name column. When a particular stock is selected, then, as shown in the lower left of Figure 7, all of the agents for that selected stock (or instrument) are listed. In a preferred embodiment of the present invention, agent factory 115 supplies a user with no more than five artificial agents 110 at any given time and for any given instrument.

For each agent listed in Figure 5, listed also is the predictability value for the agent, its state, i.e., is the agent alive (normal) or retired, and what its profit per 10 days is. As shown, there are preferably also buttons (operable via a mouse, for example) that permit a user to select individual or all agents and to request the agents' respective recommendations. Finally, an agent can be shown graphically. These latter two features are explained in more detail later herein.

Still referring to Figure 5, as already explained, the predictability value preferably is a number that varies between 0-2, 2 being the maximum predictability. Any number between 1-2 generally signifies a relatively good agent. A predictability value that is less than 0 signifies an agent that has been retired and preferably should no longer be used. Of course, the range 0-2 could be modified to any standard range, e.g., 0-10, that might be desired.

Referring still to Figure 5, the stock manager user interface screen also includes data such as the liquidity of the stock, industry sector of the stock, and other well-known parameters used to describe and/or analyze a particular stock such as earnings per annum (E12M), book value (BV), quarter when the last report was emitted (L. Rep.), price-earnings ratio (P/E), price to book value (P/BV), current price, previous price, current volume and book value to earnings (BV/E). Accordingly, this user interface screen allows a user to select a particular stock according to given fundamental/technical criteria and identify the agents associated with that stock.

The agent manager user interface screen of the present invention is depicted in Figure 6, and is similar to the stock manager, but is arranged primarily by agent rather than instrument. Accordingly, agents can more easily be selected. For instance, an agent can be listed by order of graduation date ("G Date" in Figure 6) whereby only the most recently issued agents might be selected. The agent manager user interface preferably lists all possible agents in the system, i.e. the number of instruments x 5 (assuming 5 agents are available for each instrument). In addition, the agent manager user interface screen preferably includes columns entitled price feedback indicator (PFI) and horizon (preferably listed in weeks), which indicates when the agent is expecting to find a profit-taking opportunity.

PFI is determined by presenting an agent with different hypotheses about the price during the next trading period. Each agent preferably is presented with 50 hypothetical prices regularly spaced on a logarithmic scale ranging from a 50% drop to a 50% rise from the last price. For each hypothetical future price the agent's recommendation is calculated. Of special interest here is the change between the actual recommendation at the last price and the hypothetical recommendation that the same agent would give in each future price scenario: for example an agent that is recommending to "buy" at the current price may decide to recommend to "sell" if the price were to rise 10% or more.

resistance line. That is to say even if price rises, it is likely to reach no higher than said resistance line.

Further, listed on the agent manager user interface screen shown in Figure 6 is the state of the agent 110 and its aggressiveness. Aggressiveness, in accordance with the present invention, preferably is a number that varies between 0-100 and indicates the number of trades every hundred days, on average, that the agent recommends. Accordingly, an aggressiveness value of between 25-50 is considered relatively aggressive. In contrast, a value under 10 is considered to be a relatively conservative agent.

The next column in Figure 6 is the template on which the agent is based. As explained earlier, an agent 110 is composed two brains: an A brain and a B brain. The A brain is based on one of a plurality of technical templates that seek structure in the price series. The templates include: buy and hold, moving averages, MACDs, momentum, stochastics, oscillators, composite moving averages, fourier analysis, candlesticks and the like, all with different characteristic time scales. It is noted that some of the templates look for tendencies in a price series (generically linear trends) and some look for oversold or overbought positions. For example, the moving average could be 3 day, 5 day, 10 day, 30 day, etc. The above templates are well known strategies to any practitioner in technical analysis. However, in contrast to conventional technical analysis, the present invention provides at least two important differences. i) As opposed to strategies based on fixed heuristics, e.g., buy when price crosses the moving average from below, the present invention can take the contrary strategy, i.e., sell when price crosses the moving average from below. Thus, for every conventional strategy there is an inverse strategy considered by the agents of the present invention, and ii) strategies preferably are evaluated strictly according to their predictability.

The important point, as discussed previously, is that the template that maximizes the predictability function is chosen at agent factory 115 and passed to

recommendation, and a Q factor, which is an indicator of the strength of the recommendations. Preferably, three stars indicates the strongest recommendation, two stars medium and one star indicates a fairly weak recommendation. The Q factor preferably takes into account the characteristics of the two complimentary agents that decide to open a position, such as their prediction horizons and PFI, as well as the product of their predictability values. For example three stars preferably are assigned to positions opened by two agents whose multiplied predictability exceeds 0.5, where one agent has a prediction horizon less than 5 trading periods and a positive PFI whilst the other agent has a prediction horizon of 10-20 periods and a negative PFI. A back-testing procedure using historical data is preferably used to determine which type of two-agent consensus should get a one, two or three star rating.

Also listed in Figure 7 is the price at which the position was opened and the date on which it was opened. The agents themselves are then listed so that a user can thereafter look at the CVs of either or both of these agents to learn more about them. Alternatively, the user can use the stock manager or agent manager user interface screens to learn similar information.

The next column shown in Figure 7 is "type." Agent pairs of type A have A brains that are looking for a tendency in a price (generically a linear trend) whereas agent pairs of type B have A brains that are seeking an overbought or oversold technical situation (see above discussion on technical templates). The number appended to the type letter (for example type A-2) indicates how closely the two agents' PFI and prediction horizon fit the specified type. For example, in type A positions it is preferable to have a positive PFI agent with a horizon less than five periods and a negative PFI agent with a horizon of 10-20 periods. If both agents fit this description they will open a type A-2 position. If one of the agents fits the description only approximately the type will be A-1: for example, the horizon might

be 30 periods or the PFI might be neutral (n) instead of negative (-). Type A-0 would indicate that both agents fit the description approximately.

The next column in the ticker user interface screen shown in Figure 7 is “current,” which indicates the current price of the instrument. Thus, a position is opened at a certain price and, thereafter, price changes can be monitored. The last column of this screen is a percent profit based on the price at which the position was opened and the current price.

The “closed positions” window of the ticker user interface screen preferably includes similar data as compared to the “open positions” window, but also includes the “close” date indicating when the position on the particular instrument was closed. As mentioned above, positions are opened when both agents in a complementary pair are in agreement to buy or sell. A position preferably is closed under two different circumstances: i) one of the agents changes opinion, e.g., originally both agents said sell and now one says buy; or ii) the predictability of one or both of the agents falls below a certain threshold.

Figure 8 illustrates the agent round table user interface screen. This screen provides summary-like information with respect to each instrument. In the first column, all of the instruments available in the system are listed. Then, for each instrument there is listed its price, the number of agents active for that stock and the overall predictability of the stock, i.e. the sum of the predictability values of all agents for that instrument. In addition there is a column for “PFI”. As explained above, the price feedback indicator preferably is a measure of what the expected reaction of the market will be to the recommendation of the agent. In other words, a buy could lead to a rise in the price of the bought instrument, which might lead the agent to continue to recommend a buy, which is a potentially unstable situation. A negative price feedback indicator is just the opposite. In this case, the recommendation is likely to result in equilibrium rather than instability.

roundtable user interface screen. The “force” of the stock manager’s recommendation window preferably is a number that varies between zero and a hundred, where zero is a strong sell recommendation and a hundred is a strong buy recommendation. This force value is obtained from the output of a feedforward neural network, multiplied by a factor of one hundred. Preferably, the input to the neural network is the A-brain preprocessed signal, which is described above with respect to the agent factory.

A force of 50 preferably indicates a neutral recommendation. The recommendation is “buy” if the force exceeds a threshold value that is itself greater than 50, or to “sell” when the force drops below a selling threshold. Both decision thresholds preferably are fixed, when the agent is originally created, to a value that maximizes the agent’s profits on past data.

Figure 10 illustrates exemplary graphical analyses windows for a particular instrument and a particular agent, in this case TELMEX L and TELMEXL-73, respectively. Preferably, five different windows with respective curves are simultaneously displayed. Of course, each curve by itself or other combinations of the curves may also displayed, as desired.

One curve is portfolio position, which shows the actual positions taken by the agent both during its training period in agent factory 115 and during its time with the user. A review of the agent’s CV (Figure 4) identifies the agent’s release or graduation date from the agent factory.

A force curve is also shown. The values the graph takes are the same as those encountered for Force in the agent recommendation box in the stock manager user interface screen. Generally, a force of one hundred indicates a strong buy while zero indicates a strong sell (see above).

The neural activity curve provides a measure of how many previous patterns the agent is recalling when a new price is introduced. That is, as new data is

The present invention is designed to assist decision makers with making individual correct decisions despite being faced with a complex, dynamic system such as a stock market. In this regard, agent factory 115 responds to changes in the dynamic system by continually generating new artificial agents 110 having, preferably, predictability values that are greater than a minimum threshold, e.g., a value of 1. The life span of any individual agent 110 depends significantly on the rate of change within the dynamic system. That is, in the case of a stock market that is volatile in an unpredictable way, it is likely that a typical agent's life span will be shorter as compared to a situation wherein a stock market is in a less volatile, predictable state.

In accordance with the present invention, a rental or leasing service is the preferable paradigm for agent distribution. Specifically, an investor desiring access to the agents of the present invention is preferably provided with management system 125, upon entering into a contract or leasing agreement. Management system 125 preferably resides on an electronic computer to which the investor has access either directly or over a private or public (e.g., the Internet) network.

Then, since the agents of the present invention have limited lifetimes in view of the dynamic nature of a stock market, for example, the investor preferably agrees to subscribe for a predetermined period of time (e.g., a year, or a month) to the services provided by the present invention. Specifically, once a fee has been paid and/or a contract signed, the investor becomes eligible to receive artificial agents 110 from agent factory 115. The preferred method of agent distribution is by public network such as the Internet. Those skilled in the art will appreciate that, in view of the public network transmission of the agents, encryption may be used as desired to secure their transmission to management system 125. Of course, agents may also be distributed to management system 125 via conventional storage media such as a CD ROM or even a diskette. The only requirement is that the agents arrive at management system 125 intact and uncorrupted. Agents may be updated with a

may also be implemented in agent factory 115 whereby a user is subsequently notified that a particular agent is no longer available. Management system 125 could also include the foregoing functionality whereby agent factory 115 need not be involved.

In view of all of the foregoing, it is evident that the present invention provides an improved, affordable and timely artificial agent creation, distribution and management method and system, which provides a user with a better basis for decision making.

The foregoing disclosure of embodiments of the present invention and specific examples illustrating the present invention has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many variations and modifications of the embodiments described herein will be obvious to one of ordinary skill in the art in light of the above disclosure. For example, while the agent factory and management system components have been described as separate components, they could be combined into a single system that resides in one location.

Accordingly, the scope of the invention is to be defined only by the claims appended hereto, and by their equivalents.

7. The method of claim 1, wherein step (d) is implemented when artificial agents from the first set of artificial agents are retired.

8. The method of claim 7, wherein any one artificial agent is retired when a predictability value associated therewith falls below a predetermined value.

9. The method of claim 1, wherein a price feedback indicator is associated with each artificial agent.

10. The method of claim 9, wherein the price feedback indicator is determined by querying the artificial agent as to how a recommendation would change in view of various price scenarios.

11. The method of claim 9, wherein the price feedback indicator is represented by a symbol from the group consisting of +, -, >, < and n.

12. The method of claim 1, wherein each artificial agent is pre-trained.

13. The method of claim 1, wherein steps (c) and (d) are implemented, at least in part, over an electronic network.

14. The method of claim 1, wherein each artificial agent is based on a technical template.

15. An artificial agent leasing system, comprising:
An artificial agent factory, the artificial agent factory periodically creating a pool of artificial agents;
means for distributing a first set of artificial agents from the pool of artificial agents;

means for determining when to retire any one artificial agent; and
means for distributing at least one new artificial agent from the pool of artificial agents, different from any artificial agent in the first set of artificial agents.

16. The system of claim 15, further comprising an artificial agent management system.

17. The system of claim 16, wherein the artificial agent management system comprises a graphical user interface.

18. The system of claim 16, wherein the artificial agent management system is in communication with an automated trade clearing system.

19. The system of claim 15, wherein the agent factory creates artificial agents each having a predictability value.

20. The system of claim 15, wherein each agent is self-monitoring.

21. The system of claim 15, wherein the means for distributing comprises at least one of a diskette, a CD ROM and an electronic network.

22. The system of claim 15, wherein each distributed artificial agent is pre-trained with a predetermined decision making strategy.

23. The system of claim 22, wherein the strategy results in a recommendation to one of buy, sell and hold a financial instrument.

24. A consulting system, comprising:
a means for distributing artificial agents; and

creating the artificial agents;
distributing the artificial agents;
monitoring an expected future performance of the artificial agents that have been distributed; and

if the expected future performance of one or more of the artificial agents that has been distributed has fallen below a predetermined level, creating new artificial agents and distributing the new artificial agents.

32. The method of claim 31, wherein the expected future performance is monitored at a user location.

33. The method of claim 31, wherein the request is one of a subscription and a leasing agreement.

34. The method of claim 31, wherein each artificial agent provides a financial trading recommendation.

35. The method of claim 34, wherein each artificial agent has a distinct trading strategy.

36. The method of claim 31, wherein the step of distributing comprises transmitting data over an electronic network.

37. The method of claim 36, wherein the electronic network is the Internet.

38. The method of claim 31, wherein the expected future performance of an artificial agent is associated with a predictability of a decision making strategy.

artificial agents created based, respectively, at least in part on different technical analysis templates being applied to historical price time series information; and

a management system, in communication with real time market data, operable to (i) receive the artificial agents, (ii) display characteristic information with respect to each received artificial agent and (iii) inform a user of a specific recommendation made by at least one of the artificial agents.

47. The system of claim 46, wherein the management system is operable as a tool for at least one of an individual investor, an institutional investor, a fund manager and a market maker.

48. The system of claim 46, wherein a predictability value is associated with each artificial agent.

49. The system of claim 46, further comprising an automated trade clearing system in communication with the management system.

50. The system of claim 46, wherein the specific recommendation is one of buy, sell and hold.

51. The system of claim 46, wherein a price feedback indicator is associated with each artificial agent.

52.. The system of claim 51, wherein the price feedback indicator is determined by querying the artificial agent as to how a recommendation would change in view of various price scenarios.

59. The method of claim 58, wherein the expected future performance is based on a predictability value.

60. The method of claim 58, wherein the artificial agent is self-monitoring.

61. The method of claim 58, further comprising determining the artificial agent's price feedback indicator by:

- (a) presenting the agent with different hypotheses about the price of the financial instrument during a subsequent trading period;
- (b) determining the artificial agent's recommendation for each of the different hypotheses; and
- (c) analyzing the resulting recommendations.

62. An artificial agent system, comprising:
an artificial agent management system in communication with real time data; and

a plurality of artificial agents stored in the management system, wherein each artificial agent is associated with a predictability value and wherein each artificial agent is self-monitoring, whereby any artificial agent of the plurality of agents stored in the management system is capable of retiring itself from further decision making.

63. The system of claim 62, wherein the predictability value changes in view of the real time data.

64. The system of claim 62, further comprising user interface screens for monitoring the plurality of artificial agents.

ABSTRACT

A system and method for creating, distributing and managing artificial agents suited particularly to the financial services industry. Artificial agents are created based on a combination of (i) a technical analysis template as it is applied to historic price series data of a financial instrument and (ii) neural network intelligence. A predictability function is preferably assigned to each artificial agent during its creation enabling the agent to monitor the evolution of its environment and anticipate its own future performance. The artificial agents are distributed, preferably under a subscription or leasing service arrangement, to a user's management system, which provides a graphical user interface to permit a user to review, order, select and enable selected artificial agents. Once enabled, the artificial agents provide specific trading recommendations, including buy, sell and hold, and an estimate of their future performance. Artificial agents are retired periodically as their respective trading strategies become obsolete, i.e. their value of predictability falls below a predetermined level.

1683323

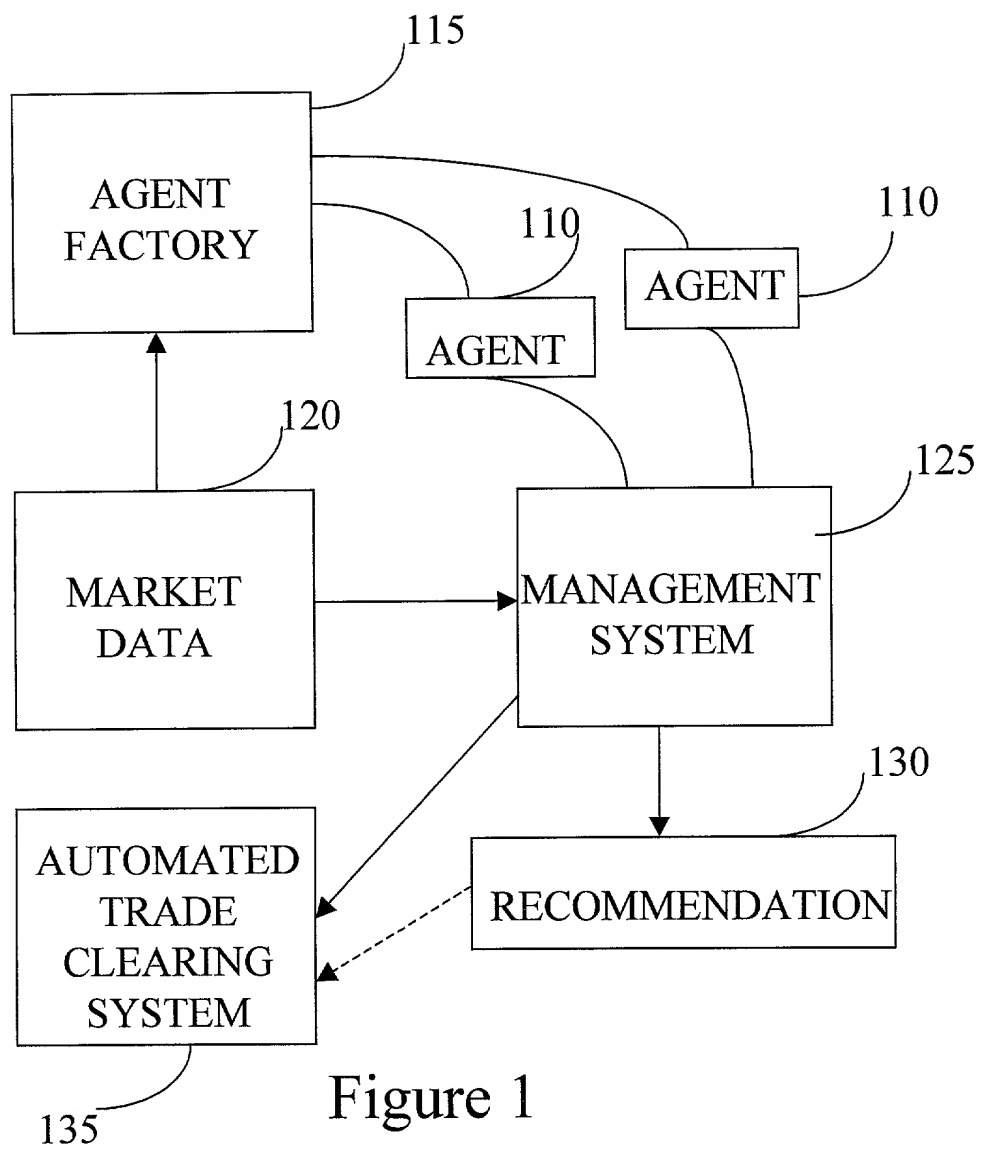


Figure 1

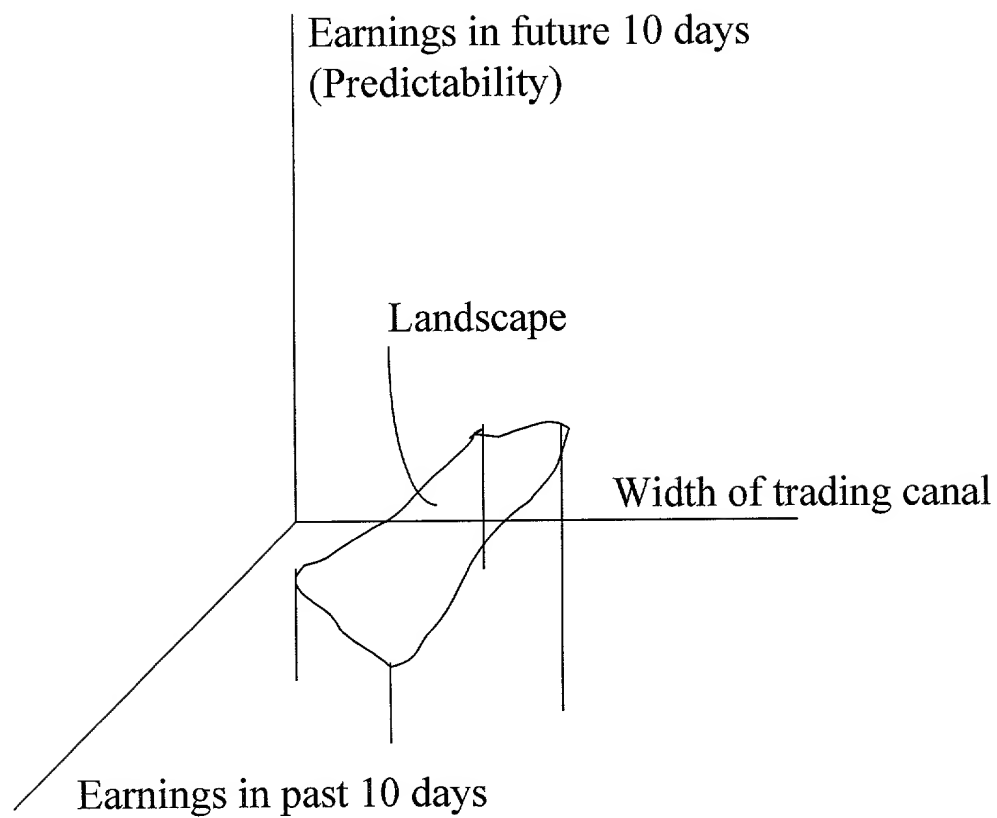



FIGURE 3



Docket No.
ADT0001-US

Declaration and Power of Attorney For Patent Application

English Language Declaration

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

"SYSTEM AND METHOD FOR CREATING, DISTRIBUTING AND MANAGING ARTIFICIAL AGENTS"

the specification of which

(check one)

☒ is attached hereto.

☐ was filed on _____ as United States Application No. or PCT International Application Number _____ and was amended on _____ (if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b) of any foreign application(s) for patent or inventor's certificate, or Section 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate or PCT International application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)			Priority Not Claimed
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>

I hereby claim the benefit under 35 U.S.C. Section 119(e) of any United States provisional application(s) listed below:

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

I hereby claim the benefit under 35 U. S. C. Section 120 of any United States application(s), or Section 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. Section 112, I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, C. F. R., Section 1.56 which became available between the filing date of the prior application and the national or PCT International filing date of this application:

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.